

Why does Monte Carlo Fail to Work Properly in High-Dimensional Optimization Problems?

Boris Polyak, Pavel Shcherbakov

Institute for Control Science, RAS, Moscow, Russia

E-mails: boris@ipu.ru; cavour118@mail.ru

July 5, 2016

Abstract

The paper proposes an answer to the question formulated in the title.

1 Introduction

After the invention of the Monte Carlo (MC) paradigm by S. Ulam in the late 1940s, it has become extremely popular in numerous application areas such as physics, biology, economics, social sciences, and other areas. As far as mathematics is concerned, Monte Carlo methods showed themselves exceptionally efficient in the simulation of various probability distributions, numerical integration, estimation of the mean values of the parameters, etc. [1, 2]. The salient feature of this approach to solution of various problems of this sort is that “often,” it is dimension-free in the sense that, given N samples, the accuracy of the result does not depend on the dimension of the problem.

On the other hand, applications of the MC paradigm in the area of optimization are not that triumphant. In this regard, problems of global optimization deserve special attention. As explained in [3] (see beginning of Chapter 1.2), “*In global optimization, randomness can appear in several ways. The main three are: (i) the evaluations of the objective function are corrupted by random errors; (ii) the points x_i are chosen on the base of random rules, and (iii) the assumptions about the objective function are probabilistic.*” Pertinent to the exposition of this paper is only case (ii). Monte Carlo is the simplest, brute force example of randomness-based methods (in [3] it is referred to as “Pure Random Search”). With this method, one samples points uniformly in the feasible domain, computes the values of the objective function, and picks the record value as the output.

Of course, there are dozens of more sophisticated stochastic methods such as simulated annealing, genetic algorithms, evolutionary algorithms, etc.; e.g., see [3, 4, 5, 6, 7] for an incomplete list of relevant references. However most of these methods are heuristic in nature; often, they lack rigorous justification, and the computational efficiency is questionable. Moreover, there exist pessimistic results on “insolvability of global optimization problems.” This phenomenon has first been observed as early as in the monograph [8] by A. Nemirovskii and D. Yudin, both in the deterministic and stochastic optimization setups (see Theorem,

Section 1.6 in [8]). Specifically, the authors of [8] considered the minimax approach to the minimization of the class of Lipschitz functions and proved that, no matter what is the optimization method, it is possible to construct a problem which will require exponential (in the dimension) number of function evaluations. The “same” number of samples is required for the simplest MC method. Similar results can be found in [9], Theorem 1.1.2, where the construction of “bad” problems is exhibited. Below we present another example of such problems (with very simple objective functions, close to linear ones) which are very hard to optimize. Concluding this brief survey, we see that any advanced method of global optimization cannot outperform Monte Carlo when optimizing “bad” functions.

This explains our interest in the MC approach as applied to the optimization setup. In spite of the pessimistic results above, there might be a belief that, if Monte Carlo is applied to a “good” optimization problem (e.g., a convex one), the results would not be so disastrous. The goal of the present paper is to clarify the situation. We examine the “best” optimization problems (the minimization of a linear function on a ball or on a cube) and estimate the accuracy of the Monte Carlo method. Unfortunately, the dependence on the dimension remains exponential, and practical solution of these simplest problems via such an approach is impossible for high dimensions.

The second goal of the paper is to exhibit the same situation with multiobjective optimization [10]. We treat methods for the Pareto set discovery via the Monte Carlo technique and estimate their accuracy, which happens to be poor for large dimensions. These results are instructive for multiobjective optimization, because there exist many methods based on a similar approach (with regular grids instead of random grids), see [11, 12].

An intuitive explanation of the effects under consideration can rely on the geometric nature of multidimensional spaces. Numerous facts of this sort can be found in Chapter 2 of the book [13], which is available in the internet. The titles of sections in [13] are very impressive: “The Volume is Near the Equator,” “The Volume is in a Narrow Annulus,” “The Surface Area is Near the Equator.” Some of the results in the present paper clarify these statements by providing rigorous closed-form estimates for the minimum number of random points in the ball-shaped sets required to assess, with a given probability, the optimum of a linear function with given accuracy. These estimates are based on our previous results on the properties of the uniform distribution over a ball, [14] (see Section 2.2).

As far as the geometry of many-dimensional spaces is concerned, the highly advanced monograph [15] is worth mentioning; it deals with the geometrical structure of finite dimensional normed spaces, as the dimension grows to infinity, and presents numerous deep mathematical results in the area.

The rest of the paper is organized as follows. In Section 2, we propose a motivating example, formulate the two optimization problems, scalar and multiobjective, considered within the Monte Carlo setup, and present two known theorems on the uniform distribution

over the l_2 -norm ball. These theorems will be used in Section 3 to derive new results related to the two optimization problems of interest for the case of the l_2 -norm ball. Section 4 deals with the scalar optimization problem for the case where Q is a box; use of various deterministic grids are also discussed. Brief concluding remarks are given in the last section.

A preliminary version of this paper is [16]. Several refinements are performed in the current text; first, we changed the overall structure of the exposition; then, we provide a new result on the probability of the empirical maximum of a linear function on a ball (Section 3.1), next, we add a result on the expected value (end of Section 3.2), present closed-form results for the l_∞ - and l_1 -norm balls (Section 4.1), discuss deterministic grids over a box (Section 4.2), and accompany these new results with numerical illustrations. Finally, the introduction section and the bibliography list are considerably modified and extended and various typos and inaccuracies are corrected.

2 Statement of the Problem

In this section, we propose a motivation for the research performed in this paper, formulate the problems of interest, and present two known facts which form the basis for deriving the new results in Section 3.

2.1 A Motivating Example

To motivate our interest in the subject of this paper, we present a simple example showing failure of stochastic global optimization methods in high-dimensional spaces. This example is constructed along the lines suggested in [8] (also, see [9], Theorem 1.1.2) and is closely related to one of the central problems discussed below, the minimization of a linear function over a ball in \mathbb{R}^n .

Consider an unknown vector $c \in \mathbb{R}^n$, $\|c\| = 1$, and the function

$$f(x) = \min \left\{ 99 - c^\top x, (c^\top x - 99)/398 \right\}$$

to be minimized over the Euclidean ball $Q \subset \mathbb{R}^n$ of radius $r = 100$ centered at the origin. Obviously, the function has one local minimum $x_1 = -100c$, with the function value $f_1 = -0.5$, and one global minimum $x^* = 100c$, with the function value $f^* = -1$. The objective function is Lipschitz with Lipschitz constant equal to 1, and $\max f(x) - \min f(x) = 1$.

Any standard (not problem-oriented) version of stochastic global search (such as multi-start, simulated annealing, etc.) will miss the domain of attraction of the global minimum with probability $1 - V^1/V^0$, where V^0 is the volume of the ball Q , and V^1 is the volume of the set $C = \{x \in Q: c^\top x \geq 99\}$. In other words, the probability of success is equal to

$$P = \frac{V^1}{V^0} = \frac{1}{2} I \left(\frac{2rh - h^2}{r^2}; \frac{n+1}{2}, \frac{1}{2} \right),$$

where $I(x; a, b)$ is the regularized incomplete beta function with parameters a and b (for use of this function, also see Theorem 1 in Section 3), and h is the height of the spherical cap C ; in this example, $h = 1$. This probability quickly goes to zero as the dimension of the problem grows; say, for $n = 15$ it is of the order of 10^{-15} . Hence, any “advanced” method of global optimization will find the minimum with relative error not less 50%; moreover, such methods are clearly seen to be no better than a straightforward Monte Carlo sampling. The same is true if our goal is to estimate the minimal value of the function f^* (not the minimum point x^*). Various methods based on ordered statistics of sample values (see Section 2.3 in [3]) fail to reach the set C with high probability, so that the prediction will be close to $f_1 = -0.5$ instead of $f^* = -1$.

These observations motivate our interest in the analysis of the performance of the MC schemes in optimization.

2.2 The Two Optimization Settings

Let $Q \subset \mathbb{R}^n$ denote a unit ball in one or another norm and let $\xi^{(i)}|_1^N = \{\xi^{(1)}, \dots, \xi^{(N)}\}$ be a multisample of size N from the uniform distribution $\xi \sim \mathcal{U}(Q)$.

We are targeted at solving the problems of the following sort.

I. Scalar Optimization: Given the scalar-valued linear function

$$g(x) = c^\top x, \quad c \in \mathbb{R}^n, \quad (1)$$

defined on the unit ball $Q \subset \mathbb{R}^n$, estimate its maximum value from the multisample.

More specifically, let η^* be the true maximum of $g(x)$ on Q and let

$$\eta = \max\{g^{(1)}, \dots, g^{(N)}\}, \quad g^{(i)} = g(\xi^{(i)}), \quad i = 1, \dots, N, \quad (2)$$

be the empirical maximum; we say that η approximates η^* *with accuracy at least δ* if

$$\frac{\eta^* - \eta}{\eta^*} \leq \delta.$$

Then the problem is: *Given a probability level $p \in (0, 1)$ and accuracy $\delta \in (0, 1)$, determine the minimal length N_{\min} of the multisample such that, with probability at least p , the accuracy of approximation is at least δ (i.e., with high probability, the empirical maximum nicely evaluates the true one).*

These problems are the subject of discussion in Sections 3.1 and 4.1.

II. Multiobjective Optimization: Consider now $1 < m < n$ scalar functions $g_j(x)$, $j = 1, \dots, m$, and the image of Q under these mappings. The problem is to “characterize” the boundary of the image set $g(Q) \subset \mathbb{R}^m$ via the multisample $\xi^{(i)}|_1^N$ from Q .

In rough terms, the problem is: *Determine the minimal sample size N_{\min} which guarantees, with high probability, that the image of at least one sample fall close to the boundary of $g(Q)$.*

For the case where Q is the Euclidean ball, the mappings $g_j(x)$ are linear, and $m = 2$, this problem is discussed in Section 3.2; various statistics (such as the cumulative distribution function, mathematical expectation, mode) of a specific random variable associated with image points are evaluated.

2.3 Supporting Material

The results presented in Section 3 below are based on the following two facts established in [14]; they relate to the probability distribution of a specific linear or quadratic function of the random vector uniformly distributed on the Euclidean ball.

Fact 1 [14]. *Let the random vector $\xi \in \mathbb{R}^n$ be uniformly distributed on the unit Euclidean ball $Q \subset \mathbb{R}^n$. Assume that a matrix $A \in \mathbb{R}^{m \times n}$ has rank $m \leq n$. Then the random variable*

$$\rho \doteq \left((AA^\top)^{-1} A\xi, A\xi \right)$$

has the beta distribution $\mathcal{B}(\frac{m}{2}, \frac{n-m}{2} + 1)$ with probability density function

$$f_\rho(x) = \begin{cases} \frac{\Gamma(\frac{n}{2} + 1)}{\Gamma(\frac{m}{2})\Gamma(\frac{n-m}{2} + 1)} x^{\frac{m}{2}-1} (1-x)^{\frac{n-m}{2}} & \text{for } 0 \leq x \leq 1, \\ 0 & \text{otherwise,} \end{cases} \quad (3)$$

where $\Gamma(\cdot)$ is the Euler gamma function.

Alternatively, the numerical coefficient in (3) writes $\frac{\Gamma(\frac{n}{2}+1)}{\Gamma(\frac{m}{2})\Gamma(\frac{n-m}{2}+1)} = 1/B(\frac{m}{2}, \frac{n-m}{2} + 1)$, where $B(\cdot, \cdot)$ is the beta function.

The second fact is an asymptotic counterpart of Fact 1.

Fact 2 [14]. *Assume that for every $n \geq m$, the matrix $A_{(n)} \in \mathbb{R}^{m \times n}$ has rank m , and $\xi_{(n)}$ is a random vector uniformly distributed over the unit ball Q in \mathbb{R}^n . Then, as $n \rightarrow \infty$, the random vector*

$$\rho_{(n)} = n^{1/2} (A_{(n)} A_{(n)}^\top)^{-1/2} A_{(n)} \xi_{(n)}$$

tends in distribution to the standard Gaussian vector $\mathcal{N}(0, \mathbf{I}_m)$, where \mathbf{I}_m is the identity $m \times m$ -matrix.

Note that for n fixed, we have

$$\|\rho_{(n)}\|^2 = n\rho; \quad (4)$$

i.e., Facts 2 and 1 characterize the asymptotic distribution of the vector $\rho_{(n)}$ and exact distribution of its squared norm (normalized by the dimension).

3 Main Results: Ball-Shaped Sets

In this section we analyse the two optimization settings formulated in Section 2.2 for Q being the n -dimensional unit l_2 -ball.

3.1 Scalar Optimization

We consider the scalar case (1) and discuss first a qualitative result that follows immediately from Fact 1.

Without loss of generality, let $c = (1, 0, \dots, 0)^\top$, so that the function $g(x) = x_1$ takes its values on the segment $[-1, 1]$, and the true maximum of $g(x)$ on Q is equal to 1 (respectively, -1 for the minimum) and is attained with $x = c$ (respectively, $x = -c$). Let us compose the random variable

$$\rho = g^2(\xi),$$

which is the squared first component ξ_1 of ξ . By Fact 1 with $m = 1$ (i.e., $A = c^\top$), for the probability density function (pdf) of ρ we have

$$f_\rho(x) = \frac{\Gamma(\frac{n}{2} + 1)}{\Gamma(\frac{1}{2})\Gamma(\frac{n+1}{2})} x^{-\frac{1}{2}}(1-x)^{\frac{n-1}{2}} \doteq \beta_n x^{-\frac{1}{2}}(1-x)^{\frac{n-1}{2}}. \quad (5)$$

Straightforward analysis of this function shows that, as dimension grows, the mass of the distribution tends to concentrate closer the origin, meaning that the random variable (r.v.) ρ is likely to take values which are far from the maximum, equal to unity. To illustrate, Fig. 1 depicts the plot of the pdf (5) for $n = 20$.




Figure 1: The probability density function (5) for $n = 20$.

We next formulate the following rigorous result.

Theorem 1. *Let ξ be a random vector uniformly distributed over the unit Euclidean ball Q and let $g(x) = x_1$, $x \in Q$. Given $p \in (0, 1)$ and $\delta \in (0, 1)$, the minimal sample size N_{\min}*

that guarantees, with probability at least p , for the empirical maximum of $g(x)$ to be at least a δ -accurate estimate of the true maximum, is given by

$$N_{\min} = \frac{\ln(1-p)}{\ln\left[\frac{1}{2} + \frac{1}{2}I\left((1-\delta)^2; \frac{1}{2}, \frac{n+1}{2}\right)\right]}, \quad (6)$$

where $I(x; a, b)$ is the regularized incomplete beta function with parameters a and b .

Clearly, a correct notation should be $N_{\min} = \lceil \cdot \rceil$, i.e., rounding toward the next integer; we omit it, but it is implied everywhere in the sequel.

Proof We specify sample size N , and let $\xi^{(i)}|_1^N$ be a multisample from the uniform distribution on Q ; also introduce the random variable

$$\eta = \max_{1 \leq i \leq N} g(\xi^{(i)}), \quad (7)$$

the empirical maximum of the function $g(x) = a^\top x$, $a = (1, 0, \dots, 0)^\top$, from this multisample. We now estimate the probability $\mathbf{P}\{\eta > 1 - \delta\}$.

By Fact 1, the pdf of the r.v. $\rho = g^2(\xi)$ is given by (5), and its cumulative distribution function (cdf) is known to be referred to as the *regularized incomplete beta function* $I(x; \frac{1}{2}, \frac{n+1}{2})$ with parameters $\frac{1}{2}$ and $\frac{n+1}{2}$, [17]. Due to the symmetry of the distribution of ξ_1 , we have $\mathbf{P}\{\rho > (1 - \delta)^2\} = 2\mathbf{P}\{\xi_1 > 1 - \delta\}$, so that $\mathbf{P}\{\xi_1 > 1 - \delta\} = \frac{1}{2} - \frac{1}{2}I((1 - \delta)^2; \frac{1}{2}, \frac{n+1}{2})$. Respectively, $\mathbf{P}\{\xi_1 \leq 1 - \delta\} = \frac{1}{2} + \frac{1}{2}I((1 - \delta)^2; \frac{1}{2}, \frac{n+1}{2})$ and $\mathbf{P}\{\eta \leq 1 - \delta\} = \left[\frac{1}{2} + \frac{1}{2}I((1 - \delta)^2; \frac{1}{2}, \frac{n+1}{2})\right]^N$, so that finally

$$\mathbf{P}\{\eta > 1 - \delta\} = 1 - \left[\frac{1}{2} + \frac{1}{2}I((1 - \delta)^2; \frac{1}{2}, \frac{n+1}{2})\right]^N. \quad (8)$$

Letting $\mathbf{P}\{\eta > 1 - \delta\} = p$ and inverting the last relation, we arrive at (6). \square

Numerical values of the function $I(x; a, b)$ can be computed via use of the MATLAB routine `betainc`. For instance, with modest values $n = 10$, $\delta = 0.05$, and $p = 0.95$, this gives $N_{\min} \approx 8.9 \cdot 10^6$, and this quantity grows quickly as the dimension n increases.

Since we are interested in small values of δ , i.e., in x close to unity, a “closed-form” lower bound for N_{\min} can be computed as formulated below.

Corollary 1. *In the conditions of Theorem 1*

$$N_{\min} > N_{\text{appr}} = \frac{\ln(1-p)}{\ln\left[1 - \frac{\beta_n}{n+1} \frac{1}{1-\delta} (2\delta - \delta^2)^{(n+1)/2}\right]},$$

where $\beta_n = \frac{\Gamma(\frac{n}{2}+1)}{\Gamma(\frac{1}{2})\Gamma(\frac{n+1}{2})} = 1/B(\frac{1}{2}, \frac{n+1}{2})$.

Proof We have

$$\begin{aligned}
I(x; \tfrac{1}{2}, \tfrac{n+1}{2}) &= \beta_n \int_0^x t^{-1/2} (1-t)^{(n-1)/2} dt \\
&= \beta_n \int_0^1 t^{-1/2} (1-t)^{(n-1)/2} dt - \gamma_n \int_x^1 t^{-1/2} (1-t)^{(n-1)/2} dt \\
&> 1 - \beta_n \int_x^1 x^{-1/2} (1-t)^{(n-1)/2} dt \quad [\text{since } t^{-1/2} < x^{-1/2} \text{ for } x < t < 1] \\
&= 1 - \beta_n x^{-1/2} \int_0^{1-x} v^{(n-1)/2} dv \quad [v = 1-t] \\
&= 1 - \beta_n \frac{2}{n+1} x^{-1/2} (1-x)^{(n+1)/2},
\end{aligned}$$

so that from (8) we obtain

$$P\{\eta > 1 - \delta\} > 1 - \left[1 - \frac{\beta_n}{n+1} \frac{1}{1-\delta} (2\delta - \delta^2)^{(n+1)/2}\right]^N$$

and

$$N_{\text{appr}} = \frac{\ln(1-p)}{\ln\left[1 - \frac{\beta_n}{n+1} \frac{1}{1-\delta} (2\delta - \delta^2)^{(n+1)/2}\right]} < N_{\min}.$$

Proof is complete. □

Further simplification of the lower bound can be obtained:

$$N_{\text{appr}} > \tilde{N}_{\text{appr}} = -\frac{\ln(1-p)}{\sqrt{2\pi(n+1)} \frac{1}{1-\delta} (2\delta - \delta^2)^{(n+1)/2}}.$$

This is doable by noting that $\ln(1-\varepsilon) \approx -\varepsilon$ for small $\varepsilon > 0$ and using the approximation $B(a, b) \approx \Gamma(a)b^{-a}$ for the beta function with large b . These lower bounds are quite accurate; for instance, with $n = 10$, $\delta = 0.05$, and $p = 0.95$, we have $N_{\min} \approx 8.8694 \cdot 10^6$, while $N_{\text{appr}} \approx 8.7972 \cdot 10^6$ and $\tilde{N}_{\text{appr}} = 8.5998 \cdot 10^6$.

3.2 Multiobjective Optimization

Consider a (possibly nonlinear) mapping $g : \mathbb{R}^n \rightarrow \mathbb{R}^m$, $n \gg m > 1$; the goal is to characterize the boundary of the image of a set $Q \subset \mathbb{R}^n$ under the mapping g . Apart from being of independent interest, this problem emerges in numerous applications. In particular, if a special part of the boundary, the *Pareto front* [10] is of interest, we arrive at a multiobjective optimization problem. Numerous examples (e.g., see [18]) show that, for n large, the images of the points sampled randomly uniformly in Q may happen to fall deep inside the true image set, giving no reasonable description of the boundary and the Pareto front of $g(Q)$.

We first present a qualitative explanation of this phenomenon by using the setup of Fact 2; i.e., the set Q is the unit Euclidean ball and the mappings are linear.

Since the squared norm of a standard Gaussian vector in \mathbb{R}^m has the χ^2 -distribution $\mathcal{C}(m)$ with m degrees of freedom [17], from Fact 2 and (4) we obtain

$$n\rho \rightarrow \mathcal{C}(m)$$

in distribution as $n \rightarrow \infty$. This is in compliance with the well-known result in the probability theory, namely, $\nu_2 \mathcal{B}(\nu_1, \nu_2) \rightarrow \mathcal{C}(\nu_1)$ in distribution as $\nu_2 \rightarrow \infty$, [17]; here, $\mathcal{B}(\nu_1, \nu_2)$ stands for the r.v. having the beta distribution with shape parameters ν_1, ν_2 . For $\nu_1 = 1$ (i.e., $m = 2$, the case most relevant to applications), Fig. 2 depicts the plots of the cumulative distribution functions $\mathcal{B}(\nu_1, \nu_2)$ (see (11) below for the explicit formula) for $\nu_2 = 1, 2, 5, 10, 20, 40$ (i.e., $n = 2, 4, 10, 20, 40, 80$).




Figure 2: Cumulative distribution functions $\mathcal{B}(1, \nu_2)$ for $\nu_2 = 1, 2, 5, 10, 20, 40$.

Hence, Fact 2 immediately implies the following important conclusion: Linear transformations essentially change the nature of the uniform distribution on a ball. Namely, as the dimension of the vector ξ grows, with the rank of the transformation matrix A being unaltered, the distribution of the vector $A\xi$ tends to “concentrate closer the center” of the image set.

We now turn to Fact 1 and provide quantitative estimates; to this end, consider the simple case where $m = 2$ and the two mappings are linear:

$$g_1(x) = c_1^\top x, \quad \|c_1\| = 1, \quad g_2(x) = c_2^\top x, \quad \|c_2\| = 1, \quad c_1^\top c_2 = 0 \quad (9)$$

(i.e., $A = \begin{pmatrix} c_1^\top \\ c_2^\top \end{pmatrix}$ in the notation of Fact 1); for instance, c_1, c_2 may be any two different unit coordinate vectors, so that $g_1(x) = x_i$ and $g_2(x) = x_j$, $i \neq j$, are the two different components of x . Then the image of Q is the unit circle centered at the origin.

Introduce now the random variable

$$\rho = g_1^2(\xi) + g_2^2(\xi), \quad (10)$$

the squared norm of the image of $\xi \sim \mathcal{U}(Q)$ under mapping (9) (i.e., $\rho = \xi_i^2 + \xi_j^2$). Then, by Fact 1 with $m = 2$, we have the closed-form expressions for the cdf F_ρ and pdf f_ρ of the r.v. ρ :

$$F_\rho(x) = \begin{cases} 0 & \text{for } x < 0, \\ 1 - (1 - x)^{\frac{n}{2}} & \text{for } 0 \leq x \leq 1, \\ 1 & \text{for } x > 1; \end{cases} \quad (11)$$

$$f_\rho(x) = \begin{cases} \frac{n}{2}(1 - x)^{\frac{n}{2}-1} & \text{for } 0 < x < 1, \\ 0 & \text{otherwise.} \end{cases} \quad (12)$$

With these in mind, let us evaluate the minimal length N of the multisample that guarantees a given accuracy with a given probability. To this end, recall that, given a multisample $\zeta^{(i)}|_1^N$ from the scalar cdf $F_\zeta(x)$ with pdf $f_\zeta(x)$, the random variable

$$\eta = \max\{\zeta^{(1)}, \dots, \zeta^{(N)}\}$$

has the cumulative distribution function $F_\eta(x) = F_\zeta^N(x)$ with pdf

$$f_\eta(x) = F'_\eta(x) = N f_\zeta(x) F_\zeta^{N-1}(x),$$

which is, in our case (11)–(12) writes

$$f_\eta(x) = \frac{Nn}{2}(1 - x)^{n/2-1} \left(1 - (1 - x)^{n/2}\right)^{N-1}. \quad (13)$$

We next evaluate several statistics of the r.v. $\eta = \max\{\rho^{(1)}, \dots, \rho^{(N)}\}$.

Probability: The theorem below determines the minimal sample size N_{\min} that guarantees, with high probability, that a random vector $\xi \sim \mathcal{U}(Q)$ be mapped close to the boundary of the image set.

Theorem 2. *Letting ξ be the random vector uniformly distributed over the unit Euclidean ball $Q \subset \mathbb{R}^n$, consider the linear mapping $g(\cdot)$ as in (9). Given $p \in (0, 1)$ and $\delta \in (0, 1)$, the minimal sample size N_{\min} that guarantees, with probability at least p , that at least one sample be mapped at least δ -close to the boundary of the image set, is given by*

$$N_{\min} = \frac{\ln(1 - p)}{\ln\left(1 - (2\delta - \delta^2)^{n/2}\right)}.$$

For small δ we have

$$N_{\min} \approx -\frac{\ln(1 - p)}{(2\delta - \delta^2)^{n/2}}.$$

Proof Let us specify sample size N and estimate the probability for a sample to be mapped close to the boundary of the image set. To this end, denote the image of $\xi \sim \mathcal{U}(Q)$ under mapping (9) by $g(\xi) = (g_1(\xi), g_2(\xi))^T$ and introduce the r.v.

$$\varkappa = \max\{\|g^{(1)}\|, \dots, \|g^{(N)}\|\},$$

the maximum of $\|g(\xi)\|$ over the multisample $\xi^{(i)}|_1^N$. Also, consider the r.v. $\rho = \rho(\xi)$ (10) for which we have

$$F_\rho(x^2) = \mathbb{P}\{\rho \leq x^2\} = \mathbb{P}\{\|g(\xi)\|^2 \leq x^2\} = 1 - (1 - x^2)^{\frac{n}{2}}$$

and the r.v. $\eta = \max\{\rho^{(1)}, \dots, \rho^{(N)}\}$, the maximum of ρ over the multisample $\xi^{(i)}|_1^N$, for which we have

$$\mathbb{P}\{\eta \leq x^2\} = F_\rho^N(x^2) = \left(1 - (1 - x^2)^{\frac{n}{2}}\right)^N. \quad (14)$$

Hence, noting that $\eta = \varkappa^2$, for a small $\delta > 0$ (i.e., letting $x = 1 - \delta$), we see that the probability for at least one sample $\xi^{(i)}$ to be mapped at least δ -close to the boundary is equal to

$$\mathbb{P}\{\varkappa > 1 - \delta\} = 1 - \left(1 - (2\delta - \delta^2)^{\frac{n}{2}}\right)^N.$$

Let $p \in (0, 1)$ be a desired confidence level; then, letting $\mathbb{P}\{\varkappa > 1 - \delta\} = p$ and inverting the last relation, we obtain the minimal required length of the multisample. The simple approximation for N_{\min} follows from the fact that $\ln(1 - \varepsilon) \approx -\varepsilon$ for small $\varepsilon > 0$. \square

To illustrate, for modest dimension $n = 10$, accuracy $\delta = 0.05$, and probability $p = 0.95$, one has to generate approximately $N = 3.4 \cdot 10^5$ random samples to obtain, with probability 95%, a point which is δ -close to the boundary of the image set. A sharper illustration of this phenomenon for $n = 50$ is given in Fig. 3, which depicts the images of $N = 100,000$ samples of $\xi \sim \mathcal{U}(Q)$ under mapping (9). None of them falls closer than 0.35 to the boundary of the image set.




Figure 3: The 2D image of the 50-dimensional ball and the result of the Monte Carlo sampling

Qualitatively, such a behavior can be explained by using geometric considerations and simple projection-type arguments discussed in [13].

Mode: The pdf (13) can be shown to be unimodal, and we find its mode by straightforward differentiating. Letting $z = 1 - x$, for the pdf we have

$$f(z) = \frac{Nn}{2} z^{n/2-1} \left(1 - z^{n/2}\right)^{N-1}.$$

Then $f'(z) = 0$ writes

$$(n/2 - 1)z^{n/2-2}(1 - z^{n/2})^{N-1} = z^{n/2-1}(N - 1)(1 - z^{n/2})^{N-2} \frac{n}{2} z^{n/2-1}.$$

Simplifying, we obtain

$$z^{n/2} = \frac{n - 2}{nN - 2},$$

hence,

$$x_{\max} = 1 - \left(\frac{n - 2}{nN - 2}\right)^{2/n}.$$

We thus arrive at the following result.

Theorem 3. *Letting ξ be the random vector uniformly distributed over the unit Euclidean ball $Q \subset \mathbb{R}^n$, consider the linear mapping $g(\cdot)$ as in (9) and the random variable*

$$\eta = \max_{i=1,\dots,N} \rho^{(i)},$$

the empirical maximum of the function $\rho(x) = \|g(x)\|^2$ from the multisample $\xi^{(i)}|_1^N$ of size N . The mode of the distribution of η is given by

$$x_{\max} = 1 - \left(\frac{n - 2}{nN - 2}\right)^{2/n}.$$

For large n we have an approximation

$$x_{\max} \approx 1 - \frac{1}{N^{2/n}}.$$

The quantity x_{\max} is seen to be essentially less than unity for large n , even if the sample size N is huge. This means the r.v. η takes values far from the boundary of the image. For instance, let $n = 20$; then, to “obtain a point in the 0.1 vicinity of the boundary,” one has to generate $N \approx 10^{10}$ random samples in Q . The family of the pdfs (13) is plotted in Fig. 4.

Expectation: We now estimate the mathematical expectation \mathbf{E} of the empirical maximum.




Figure 4: Probability density functions (13) for $n = 20$ and $N = 10^k$, $k = 2, \dots, 10$

Theorem 4. *Under the conditions of Theorem 3 we have*

$$\mathbb{E}(\eta) = 1 - \frac{2}{n} B\left(\frac{2}{n}, N + 1\right), \quad (15)$$

where $B(\cdot, \cdot)$ is the beta function.

Proof If a r.v. ζ is positive and defined on $D = [0, d]$, then the expectation

$$\mathbb{E}(\zeta) = \int_0^d (1 - F_\zeta(x)) dx,$$

where $F_\zeta(x)$ is the cdf of ζ . Hence, having N samples $\xi^{(i)}$ of $\xi \sim \mathcal{U}(Q)$ and the respective r.v. $\rho = \rho(\xi)$ (10) with support $[0, 1]$, for the r.v. $\eta = \max\{\rho^{(1)}, \dots, \rho^{(N)}\}$ we have

$$\mathbb{E}(\eta) = \int_0^1 (1 - F_\eta(x)) dx,$$

where $F_\eta(x)$ is given by (14). By change of variables $z = (1 - x)^{n/2}$, we arrive at (15). \square

For large n and N , numerical values of the expectation are close to those observed for the mode; this is seen from the shape of the pdf (13) depicted in Fig. 4. More formally, having the approximation $B\left(\frac{2}{n}, N + 1\right) \approx \Gamma\left(\frac{2}{n}\right)(N + 1)^{-2/n}$ for large N , from (15) we obtain

$$\mathbb{E}(\eta) \approx 1 - \frac{2}{n} \Gamma\left(\frac{2}{n}\right)(N + 1)^{-2/n} = 1 - \Gamma\left(\frac{2}{n} + 1\right)(N + 1)^{-2/n} \approx 1 - N^{-2/n}.$$

For instance, with $n = 20$ and $N = 10^9$, we have $\mathbb{E}(\eta) = 0.8802$ for the expectation and $x_{\max} = 0.8754$ for the mode.

4 Main Results: Box-Shaped Sets

In this section, we consider the scalar optimization problem, however, for box-shaped sets, i.e., not related to Facts 1 and 2. We consider the scalar setup described in Section 2.2 along with the deterministic approach based on use of regular grids.

4.1 A Direct Monte Carlo Approach

Consider the linear scalar optimization problem for the case where $Q = [-1, 1]^n$. Clearly, the results heavily depend on the vector c in the optimized function $g = c^\top x$; we consider two extreme cases.

Case 1. First, let $c = (1, 0, \dots, 0)^\top$ and consider the empirical maximum

$$\eta = \max\{g^{(1)}, \dots, g^{(N)}\},$$

where $g^{(i)}$ is the first component of the random vector $\xi \sim \mathcal{U}(Q)$. Specifying $\delta \in (0, 1)$, we obtain

$$\mathbb{P}\{\eta \leq 1 - \delta\} = (1 - \delta/2)^N.$$

This quantity is seen to be independent of the dimension (which is obvious as it is). Now, specifying a probability level $p \in [0, 1]$, we obtain that the minimal required sample size that guarantees accuracy δ with probability p is equal to

$$N_{\min} = \frac{\ln(1 - p)}{\ln(1 - \delta/2)}.$$

For instance, with $p = 0.95$ and $\delta = 0.1$, one has to generate just 59 points to obtain a 10%-accurate estimate of the maximum with probability 95%, independently of the dimension.

Case 2. Now let $c = (1, 1, \dots, 1)^\top$; i.e., the optimized function is $g(x) = \sum_i^n x_i$, so that the maximum is attained at $x = c^\top$ and is equal to $\eta^* = n$. In contrast to Case 1, Monte Carlo sampling exhibits a totally different behavior. Below, $\mathbf{Vol}(\cdot)$ stands for the volume of a set.

Theorem 5. *Letting ξ be the random vector uniformly distributed over the unit l_∞ -norm ball $Q = [-1, 1]^n$, consider the linear function $g(x) = \sum_i^n x_i$. Given $p \in (0, 1)$ and $\delta \in (0, 1)$, $\delta \leq 1/n$, the minimal sample size N_{\min} that guarantees, with probability at least p , for the empirical maximum of $g(x)$ to be at least a δ -accurate estimate of the true maximum, is given by*

$$N_{\min} = \frac{\ln(1 - p)}{\ln\left(1 - \frac{n^n \delta^n}{2^n n!}\right)}. \quad (16)$$

For small δ and large n we have

$$N_{\min} < -\frac{\sqrt{2\pi n}}{(\delta e/2)^n} \ln(1 - p). \quad (17)$$

Proof Let us specify a small $\delta \in (0, 1)$ and define

$$Q_\delta = \{x \in Q: \sum_{i=1}^n x_i \geq n(1 - \delta)\},$$

so that the maximum of $g(x)$, over $Q \setminus Q_\delta$ is equal to $n(1 - \delta)$. For $\delta \leq 1/n$, the set Q_δ is seen to be the simplex with $n + 1$ vertices at the points $v_0 = (1, \dots, 1)^\top$ and

$$v_j = (1, \dots, 1, \underbrace{1 - n\delta}_j, 1, \dots, 1), \quad j = 1, \dots, n,$$

with $\mathbf{Vol}(Q_\delta) = |\frac{1}{n!} \det(v_1 - v_0; \dots; v_n - v_0)| = \delta^n n^n / n!$. Since $\mathbf{Vol}(Q) = 2^n$, for $\xi \sim \mathcal{U}(Q)$ we have

$$\mathbf{P}\{\xi \in Q_\delta\} = \frac{\delta^n n^n}{2^n n!} \quad \text{and} \quad \mathbf{P}\{\xi \in Q \setminus Q_\delta\} = 1 - \frac{\delta^n n^n}{2^n n!},$$

so that

$$\mathbf{P}\{\eta > n(1 - \delta)\} = 1 - \left(1 - \frac{\delta^n n^n}{2^n n!}\right)^N.$$

Equating this probability to p and inverting this relation leads to (16).

The lower bound (17) follows immediately from Stirling's formula and the fact that $\ln(1 - \varepsilon) \approx -\varepsilon$ for small $\varepsilon > 0$. \square

For $n = 10$ and the same values $\delta = 0.1$ and $p = 0.95$, we obtain a huge $N_{\min} \approx 1.12 \cdot 10^{10}$. Even for $n = 2$, an “unexpectedly” large number $N_{\min} \approx 600$ of samples are to be drawn.

l_1 -norm ball: The setup of Case 2 is of the same flavor as the one where the set Q is the unit l_1 -norm ball, and the optimized function is $g(x) = c^\top x$ with $c = (1, 0, \dots, 0)^\top$. We have a result similar to those in Theorems 1 and 5.

Theorem 6. *Letting ξ be the random vector uniformly distributed over the unit l_1 -norm ball $Q = \{x \in \mathbb{R}^n: \sum_{i=1}^n |x_i| \leq 1\}$, consider the linear function $g(x) = x_1$. Given $p \in (0, 1)$ and $\delta \in (0, 1)$, the minimal sample size N_{\min} that guarantees, with probability at least p , for the empirical maximum of $g(x)$ to be at least a δ -accurate estimate of the true maximum, is given by*

$$N_{\min} = \frac{\ln(1 - p)}{\ln(1 - \frac{1}{2}\delta^n)}. \quad (18)$$

For small δ we have

$$N_{\min} \approx -\frac{\ln(1 - p)}{0.5 \delta^n}.$$

Proof The true maximum of $g(x)$ on Q is equal to unity; we specify accuracy $\delta \in (0, 1)$ and consider the set

$$Q_\delta = \{x \in Q: x_1 \geq 1 - \delta\}.$$

We then have

$$\mathbf{Vol}(Q) = \frac{2^n}{n!}, \quad \mathbf{Vol}(Q_\delta) = \frac{(2\delta)^n}{2 \cdot n!},$$

so that for $\xi \sim \mathcal{U}(Q)$ we obtain

$$\mathbb{P}\{\xi \in Q \setminus Q_\delta\} = \frac{\text{Vol}(Q \setminus Q_\delta)}{\text{Vol}(Q)} = 1 - \frac{1}{2}\delta^n,$$

and the rest of the proof is the same as that of the previous theorem. \square

To compare complexity associated with evaluating the optimum of a linear function over the l_2 -, l_∞ -, and l_1 -balls, we present a table showing the minimal required number of samples for $\delta = 0.05$, $p = 0.95$ and various dimensions, as per formulae (6), (16), and (18), respectively.

n	1	2	3	4	5	10	15
l_2	119	449	$1.6 \cdot 10^3$	$5.7 \cdot 10^3$	$2 \cdot 10^4$	$8.9 \cdot 10^6$	$3.6 \cdot 10^9$
l_∞	119	$2.4 \cdot 10^3$	$4.3 \cdot 10^4$	$7.2 \cdot 10^5$	$1.2 \cdot 10^7$	$1.1 \cdot 10^{13}$	10^{19}
l_1	119	$2.4 \cdot 10^3$	$4.8 \cdot 10^4$	$9.6 \cdot 10^5$	$1.9 \cdot 10^7$	$6.1 \cdot 10^{13}$	$2 \cdot 10^{20}$

Table 1: l_p -balls: Minimal required number of samples for $\delta = 0.05$ and $p = 0.95$.

These results are in consistence with intuition, since the l_1 -norm ball is “closer” in shape to the “worst-case” conic set, while the l_∞ -norm ball with $c = (1, 1, \dots, 1)^\top$ “takes an intermediate position” between l_2 and l_1 (obviously, closer to l_1).

4.2 Deterministic Grids

In this section we briefly discuss a natural alternative to random sampling the set Q . A belief is that use of various deterministic grids might outperform straightforward Monte Carlo. Again, we consider Q being the unit box $[-1, 1]^n$, and the scalar function to be optimized is $g(x) = \sum_{i=1}^n x_i$, so that the maximum is equal to n . We show that, even in such a simple setting, deterministic grids happen to be computationally intensive in high dimensions.

Uniform grid: Consider a positive integer $M > 1$ and the uniform mesh \mathcal{M} on Q , with cell-size $\Delta = 2/(M + 1)$; the mesh is assumed not to cover the boundary of Q . The total amount of points in the mesh is $\text{card } \mathcal{M} = M^n$ and the maximum of $g(x)$ over \mathcal{M} is equal to $g_{\mathcal{M}} = n(1 - \Delta)$. To guarantee relative accuracy δ of approximation, i.e., $g_{\mathcal{M}} = n \cdot (1 - \delta)$, one needs cell-size to be $\Delta = \delta$, hence, the overall number of mesh points is equal to

$$\text{card } \mathcal{M} = \left(\frac{2}{\delta} - 1\right)^n, \quad (19)$$

which amounts to a huge $M \approx 6.13 \cdot 10^{12}$ for modest $n = 10$ and $\delta = 0.1$. Interestingly, to obtain the same accuracy *with probability* $p = 0.99$, Monte Carlo requires “just” $N_{\min} = 1.7 \cdot 10^{10}$ samples!

Sobol sequences: Another type of grids that can be arranged over boxes are *low-discrepancy* point sets or *quasi-random* sequences [19]. In practice, they share many properties of *pseudorandom* numbers, e.g., such as those produced by the `rand` routine in MATLAB.

Among the variety of quasi-random sequences, so-called LP_τ sequences introduced by I.M. Sobol in 1967, [20] (also see [11, 12] for recent developments) are widely used in various application areas. This sophisticated mechanism heavily exploits the box shape of the set; it is much more efficient than purely deterministic uniform grids and may outperform straightforward Monte Carlo.

In the experiments, we considered the function $g(x) = \sum_{i=1}^n x_i$ defined on $Q = [-1, 1]^n$ and computed its maximum value over the points of an LP_τ sequence of length $N = 10^6$ for various dimensions; this was performed by using the MATLAB routine `sobolset`. The corresponding results are given in the row “ LP_τ ” of Table 2. The row “Monte Carlo” presents empirical maxima obtained by using Monte Carlo sampling with the same sample size N (averaged over 100 realizations), and the row “Uniform grid” relates to the uniform mesh of cardinality N .

n ; true max	2	3	4	5	10	15	20
Uniform grid	1.9960	2.9406	3.7576	4.4118	6.0000	7.5000	6.6667
LP_τ	1.9999	2.9792	3.8373	4.6844	7.9330	10.2542	10.9470
Monte Carlo	1.9974	2.9676	3.8731	4.6981	7.8473	10.0796	11.8560

Table 2: l_∞ -box: Empirical maxima for the three methods with $N = 10^6$.

It is seen that the uniform mesh exhibits a very poor relative performance as dimension grows, while the results of LP_τ and Monte Carlo approaches are much better and similar to each other. Clearly, the absolute values are very far from the true maxima equal to n , since $N = 10^6$ samples are not sufficient to obtain reasonable accuracy for dimensions $n > 5$.

Instead of computing the sample size for fixed values of the accuracy δ as in Table 1, here we fix the sample size and compute the empirical maxima. The reason for such an “inversion” is that, given δ and the specific linear function $g(x)$, it is not quite clear how to estimate the required length of the LP_τ -sequence. To overcome this problem, one might fix a reasonable value of accuracy, compute the minimal sample size N_{\min} required for Monte Carlo, run the LP_τ -approach with this length, and compare the results. However, for large dimensions n , the values of N_{\min} are huge, leading to very large computation times or even memory overflow.

The values obtained for the uniform mesh were computed by inverting relation (19) and using the actual grid with cardinality $\text{card } \mathcal{M} = \lceil N^{1/n} \rceil^n \geq N$, so the quantities presented in row “Uniform grid” of Table 2 are overly optimistic.

More importantly, the routine `sobolset` has several parameters, such as $s = \text{skip}$ (choice

of the initial point in the sequence) and $\ell = \texttt{leap}$ (selection of every ℓ th point). These play the similar role as the seed value in MATLAB pseudorandom number generators does, and may happen to be crucial for the quality of the resulting LP_τ sequence. In the experiments, we used the default values $s = \ell = 1$ (i.e., the first N points), though for different values of s, ℓ , the respective estimates of the empirical maximum may differ a lot. For instance, adopting $\ell = 133$, for $n = 5$ we obtain a much better estimate 4.8653, while taking (intentionally bad) $\ell = 128$ leads to very poor 2.8734.

Finally, note that applications of uniform grids and LP_τ sequences are limited to box-shaped sets. Sets different from boxes can in principle be embedded in tight enclosing boxes with subsequent use of rejection techniques; however, the rejection rate usually grows dramatically as the dimension increases.

5 Conclusions

The main contribution of the paper is a rigorous explanation of the reason why does a direct Monte Carlo approach show itself inefficient in high-dimensional optimization problems when estimating the maximum value of a function from a random sample in the domain of definition. First, attention was paid to linear functions and ball-shaped sets; using known results on the uniform distribution over the ball, we characterized the accuracy of the estimates obtained via a specific random variable associated with the function value. Also, a multiobjective optimization setup was discussed. The results obtained testify to a dramatic growth of computational complexity (required number of samples) as the dimension of the ball increases. Same flavor results are obtained for box-shaped sets; these also include analysis of deterministic grids.

The authors are indebted to an anonymous reviewer for the critical comments that led to a tangible progress in the presentation of the results.

References

- [1] Kroese, D.P., Taimre, T., Botev, Z.I.: Handbook of Monte Carlo Methods. John Wiley and Sons, New York (2011)
- [2] Tempo, R., Calafiore, G., Dabbene, F.: Randomized Algorithms for Analysis and Control of Uncertain Systems, with Applications. Springer, London (2013)
- [3] Zhigljavsky, A., Žilinskas, A.: Stochastic Global Optimization. Springer Science+Business Media, New York (2008)
- [4] Handbook of Global Optimization. Kluwer, Dordrecht (1995) vol. 1, Horst, R. and Pardalos, Panos M. (eds.); (2002) vol. 2, Pardalos, Panos M. and Romeijn, H.E. (eds.)

- [5] Simon, D.: Evolutionary Optimization Algorithms. Wiley, New York (2013)
- [6] Goldberg, D.: Genetic Algorithms in Search, Optimization and Machine Learning. Addison-Wesley, Reading, MA (1989)
- [7] Solis, F.J., Wets, R.J-B.: (1981). Minimization by random search techniques. *Mathematics of Operation Research*. **6**(1), 19–13 (1981)
- [8] Nemirovski, A., Yudin, D.B.: Problem Complexity and Method Efficiency in Optimization. John Wiley and Sons, New York (1983)
- [9] Nesterov, Yu.: Introductory Lectures on Convex Optimization: A Basic Course. Klüwer (2004)
- [10] Deb, K.: Multiobjective optimization. In: Burke, E.K. and Kendall, G. (eds.): Search Methodologies: Introductory Tutorials in Optimization and Decision Support Techniques, pp. 444–463. Springer Science+Business Media, New York (2014)
- [11] Statnikov, R., Matusov, J.: Multicriteria Analysis in Engineering Using the PSI Method with MOVI 1.0. Kluwer Academic Publishers, Dordrecht (2002)
- [12] Statnikov, R., Matusov, J., Statnikov, A.: Multicriteria engineering optimization problems: Statement, solution and applications. *J. Optim. Theory Appl.* **155**(2), 355–375 (2012)
- [13] Hopcroft, J., Kannan, R.: Foundations of Data Science. Available at <http://research.microsoft.com/en-US/people/kannan/book-no-solutions-aug-21-2014.pdf>
- [14] Polyak, B.T., Shcherbakov, P.S.: Random spherical uncertainty in estimation and robustness. *IEEE Transactions on Automatic Control*. **45**(11), 2145–2150 (2000)
- [15] Milman, V.D., Schechtman, G.: Asymptotic Theory of Finite Dimensional Normed Spaces. Springer-Verlag, Berlin (2001)
- [16] Polyak, B.T., Shcherbakov, P.S.: Failure of the Monte Carlo methods in optimization problems of high dimensions. In: Granichin, O.N. (ed.): Stochastic Optimization in Informatics, vol. 10, no. 1, pp. 89–100 (in Russian). St.-Petersburg State University (2014)
- [17] Wilks, S.S.: Mathematical Statistics. John Wiley and Sons, New York (1962)
- [18] Polyak, B., Shcherbakov, P., Khlebnikov, M.: Quadratic image of a ball: Towards efficient description of the boundary. In: International Conference on System Theory, Control, and Computing (ICSTCC 2014), pp. 105–110. IEEE CSS (2014)

- [19] Kuipers, L., Niederreiter, H.: Uniform Distribution of Sequences. John Wiley, New York (2005)
- [20] Sobol, I.M.: On the distribution of points in a cube and the approximate evaluation of integrals. USSR Computat. Maths. Math. Phys. **7**(4), 86–112 (1967)